

Table 1

I. Overview

During the past year, research has progressed along three fronts:

1. Development of a high-performance FORTRAN compiler for the Navier-Stokes Computer (NSC)
2. Identification of appropriate application codes
3. Study of an upgrade to the NSC architecture and hardware to produce a next-generation node.

These areas are summarized below.

II. Compiler Development for the Navier-Stokes Computer

Research continued on a generic prototype compiler to efficiently port code to the NSC. The initial target platform for the compiler was the NSC MiniNode to which FORTRAN code would be ported. Although the compiler is generic in its architecture, the focus in the research has been to create an initial version which efficiently accepts unmodified 'dusty deck' ANSI FORTRAN 77 source code and optimizes variable storage in memory to minimize reference conflicts. The compiler further takes advantage of the NSC architecture and maximizes the average number of floating-point and integer/logic processors utilized over the course of a flow simulation. The novel feature of the compiler include:

1. direct creation of dependency graphs from the unmodified source code,
2. high-level approximate modelling of various elements of the target computer architecture (in this case, the NSC), a short list of which includes:
 - (a) memory address computation unit,
 - (b) memory plane and cache architecture, size and update/replacement algorithms,
 - (c) ALS external and internal data paths, registers and execution modelling,
3. performance prediction of candidate code fragments produced by the compiler using the models from Item 1 above to provide feedback to the compiler during optimization, and
4. heuristics aimed at recognizing and efficiently implementing computational constructs frequently encountered in CFD algorithms (such as linear matrix operations, FFTs, conditional evaluations for numerical stability, etc.).

The compiler is parameterized, and is dubbed a Parameterized Memory/Processor (PMP) optimizing compiler. The basic features of a parallel computer architecture, such as the

94-29030



2

DTIC QUALITY INSPECTED 3

94 9 06 1 03

number, type, and behavior of memory, processors, registers, control stores, and their interconnections and couplings, are parameterized. This permits the study of the suitability of existing and proposed parallel computers in handling large scientific codes. It also provides a means by which architectural variations may improve performance, such as the addition of additional ports to memory for example.

Recent work has indicated that certain specific hardware elements (e.g. portions of the hardwired ALS fine-grain switching network and some of the conditional evaluation circuitry) of the NSC Mininode were being utilized less than one-percent of the simulated run time for the benchmark codes previously selected for the NSC. The conclusion was that the NSC node would profit from a removal of certain hardware units which are only seldom utilized. In essence, this would create a so-called 'reduced instruction set' for portions of the data path within the ALS pipeline. The net result is a significant decrease in node cost and/or a substantial increase in measured performance on most of the problems-of-interest to be studied on the NSC. This was addressed in a study of upgrades for a next-generation NSC Node.

III. NSC Application Codes

Various application areas for NSC simulation have been identified as significant and appropriate to demonstrate the compiler, especially with planned NSC Node upgrades (discussed in Section IV). These areas include:

1. Flows dominated by concentrated regions of longitudinal vorticity
2. Turbulent boundary-layer flows which may be electrically conducting.

These application areas represent complex flows for which flow control may be attempted. Early experiments at Princeton have demonstrated the potential of decreasing the swirl velocities in trailing wing-tip vortices through small modifications of the wing planform in the region of the tip. Simple point-vortex methods have been previously used to predict the wake flow based on particular wing-tip modifications. The NSC can be used to design wingtip modifications which attempt to minimize the vortex wake. Also, the NSC can simulate numerous candidate designs using automated optimal design methods coupled with three-dimensional vortex-element methods.

A second area of investigation involves the reduction of turbulent boundary-layer wall stress through the application of a time- and spatially-varying an electromagnetic body (Lorentz) force to fundamentally restructure the boundary layer. Experiments in the past year at Princeton have demonstrated that an array of surface-mount electrodes and magnets can be used to introduce large-scale waves of vorticity in the boundary-layer through the application of a distributed, time-varying applied electric field, and a distributed permanent magnetic field. These waves, when created with the appropriate amplitude and scale, were observed to decrease turbulence production and efficiently lower the drag in excess of 90%. Although the phenomena has been demonstrated in a water facility seeded with a weakly-conducting electrolyte at $Re_\theta \geq 3,000$, the basic physics of the observed effects has yet to be fully explained.

Dis.		
A-1		

EMTC simulations can be used to shed a great deal of light on the underlying phenomenology. It was determined that the NSC is an appropriate platform in terms of architecture, software, and planned hardware upgrades to explore such simulations. Appropriate algorithms, such as time-accurate solvers coupled to three-dimensional unstructured grids, or other three-dimensional DNS codes, can be modified to account for the Lorentz force as a simple body force in the initial simulations. The specific characteristics of the electric, current, and magnetic fields would be numerically determined off-line with existing standalone finite-element codes, and then imported to the NSC solvers as an uncoupled body force.

IV. Next-Generation Navier-Stokes Computer Node

The NSC architecture was initially conceived ten years ago in 1984, and was subsequently implemented in hardware. The last major revision to the hardware occurred in 1988 when dedicated floating-point processors were added. Subsequent to that time, a number of major improvements to processor design have been made by the integrated circuits industry. Thus, a portion of the recent research has focussed on assessing the viability of NSC architecture as a whole, and considering direct node-hardware upgrades.

The NSC MiniNode was the operational prototype hardware node that was used for most of these studies. The MiniNode represents the key building block of the NSC dynamically-reconfigurable parallel-processing supercomputer. The NSC is based on a small number of powerful nodes (eg. MiniNodes), where each node *running standalone* is comparable in sustained performance to current supercomputers. Each node has dynamically reconfigurable internal systolic arrays (arithmetic logic structures, or ALSs) connected via crossbar switch to multiple independent memory and address-generation modules (memory planes). Three kinds of ALSs can be used in the node: singlets with one floating-point unit, doublets with two floating-point units, and triplets with three. These computational and storage assets support the multiple levels of parallelism required for efficient solution of most numerical forms of the Navier Stokes Equations. The ALSs alone provide fine-grained support, while the node as a whole represents medium-grain hardware parallelism. Multiple nodes that are interconnected provide coarse-grain support for global domain decomposition. Overall, the architecture may be described as a continually reconfigurable Very Long Instruction Word (VLIW) machine.

The number of memory and ALS assets are used to categorize NSC nodes of various sizes. The convention is to specify the number of singlets, doublets, and triplets and memory planes in a node. Thus, an $x:y:z/m$ configuration has x singlets, y doublets, and z triplets with m memory planes. At present, the prototype MiniNode exists in hardware at Princeton University, and has a 0:2:2/4 configuration (i.e. two doublets, two triplets — 10 total floating-point units — and 4 memory planes). When hardware is upgraded, this configuration will be doubled to a 0:4:4/8 configuration. The performance of such a configuration is discussed below.

In the recent work, the NSC architecture was found to benefit from a 'streamlining' of certain hardware elements that define and control the data-flow path within the Node.

Slight modifications of the ALS, central crossbar switch, and nanocode program storage units are required to provide a substantial increase efficiency (fractional throughput) and to take advantage of the current generation of processor chips. A consequence of this is that the compiler port will be further simplified.

It was found that in a simple tradeoff analysis of the switching and ALS units in the Mininode, based in part on the experience gained from the compiler development, the ALS switching network can be minimized to reduce the amount of direct-connect hardware paths implemented by dedicated, fully general cross-point switching chips. Certain ALS support units, such as selected hardwired IF-THEN-ELSE units, can be eliminated along with nanocode storage and control circuitry.

The current Mininode has typically operated at average speeds of 100 Mflop at a (large-scale production) cost of \$25,000 per unit. Based on the compiler, new hardware, and streamlined architecture, the recent research has indicated that each next-generation NSC Node can readily have the following characteristics:

1. > 2 GIPS (billion instructions per second) peak and ~ 1.5 GIPS sustained
2. > 1 GFLOPS peak and ≥ 0.75 GFLOPS sustained
3. > 100 Mbytes high-speed memory
4. ~ 5 Gbytes rotating on-line R/W storage.
5. < \$25,000 per node (estimated 1994/95 memory and processor pricing)
6. PMP FORTRAN Compiler

Given the advances of chip design since the last NSC hardware upgrade, the overall volume of the next-generation Node is projected to *shrink* by 30%, with power consumption and heat dissipation similarly decreasing. The node architecture supports a maximum configuration of 16,379 interconnected nodes (1 node port is reserved for a front-end and 4 ports are reserved for high-speed access to secondary storage). It is anticipated that given the power of individual nodes, and the increased efficiency and ease-of-use of architectures which are not massively parallel, initial implementations of the NSC will involve far fewer nodes than the maximum supported by the architecture.

In summary, a 256-node NSC connected via crossbar is effectively coarse-grained for straight-forward domain and task decomposition and partitioning with the PMP compiler. The overall performance (peak) of such a system with the next-generation Nodes is projected to be:

1. > 0.5 TIPS (trillion instructions per second) peak and ~ 1.5 GIPS sustained
2. > 0.25 TFLOPS peak and ≥ 0.15 TFLOPS sustained on production codes
3. > 25 Gbyte high-speed memory
4. > 1 Tbyte rotating on-line R/W storage.

Such a configuration of the upgraded NSC with PMP compiler is will be cost-effective, especially when the true cost of porting large-scale aerodynamic and hydrodynamic production codes to the current crop of so-called 'general-purpose' massively-parallel computers is considered.